

Reassessing the market impact of cyber incidents: A bias-adjusted event study approach

Abstract

This study examines the impact of cyber incidents on shareholder value. Using data from 2012 to 2022, we measure abnormal stock returns around reported incidents, adjusting for event-induced variance and cross-correlation. Unlike prior research, we find no statistically significant market-wide abnormal returns once these adjustments are applied. However, data breaches stand out as particularly damaging, with average losses of -1.3% (USD -1.9 billion). The health sector is especially vulnerable, with average losses of -5.2% . Our results suggest that cyber risk is priced selectively by markets, with implications for portfolio risk assessment, sector allocation, and investment strategies.

JEL classification: C12, C23, C58, G14.

Keywords: event study, econometrics, cybersecurity.

1. Introduction

Following the hack of ICBC, Marcus Murray, the founder of Swedish cybersecurity firm Truesec, declared:

“This is a true shock to large banks around the world. The ICBC hack will make large banks around the globe race to improve their defences, starting today.”¹

This major event, which has cost the firm itself and the U.S. Treasury market’s clearing organization millions, is a single example of how detrimental cyber incidents can be to the economy. Thus, quantifying these costs is an essential starting point for cybersecurity investments by target organizations, as well as for cyber insurers and cybersecurity service providers to determine their market size. Given the publicly available information that stock markets provide researchers, a popular econometric approach, the event study, is the default method for systematically assessing these costs and explaining their sources of variation. Since Campbell, Gordon, Loeb, and Zhou (2003), the first to adopt this approach for cyberattacks, a significant strand of literature in the economics of information security has used it, with conflicting results at a granular level but yielding overall support for the view that cyberattacks or more broadly cyber incidents do affect firms’ value.

In this paper, we revisit the impact of cyber incidents on listed firms using observed adjusted returns on their stocks and advanced event study methods. Although such studies already exist in cybersecurity and financial literature, we are the first to adopt the following setting. First, we use newswire headlines that we filter to ensure they bring relevant information regarding cyber incidents. Next, we use state-of-the-art estimation methods to account for the cross-correlation of errors at the estimation stage of our cumulative abnormal returns (CARs) using a seemingly unrelated regression (SUR) estimation. This ensures that the reported CAR coefficients are unbiased regarding economic magnitude. We find overall negative abnormal returns of about -0.89% irrespective of the target firm or the type of cyber incident. Next, we adjust the standard errors using the Boehmer, Musumeci, and Poulsen (1991) and the Kolari and Pynnonen (2010) corrections for event-induced variance and cross-correlation, respectively. This is a crucial and significant caveat in the existing literature. First, the proximity of firms suffering from cyber incidents is strong in terms of characteristics and behavior, thereby generating implicit cross-correlation in their returns.

¹<https://time.com/6333716/china-icbc-bank-hack-usb-stick-trading/>

Second, it is likely that a cyber incident not only affects the target firm but also the entire market. These two biases would then favor an underestimation of standard errors.

Unadjusted standard errors yield CARs with significance well below the 1% level. These results translate into overall economic costs of USD 277 billion market capitalization reduction over the period (USD 25 billion per year), with a median cost of USD 123 million for the US-listed firms at each cyber incident. However, once we account for the aforementioned biases, we find no statistical significance to remain at the aggregated level.² This goes against the existing consensus in the event study literature on cyberattacks. We subsequently turn to a more granular analysis of these effects, adding variables controlling for the type of cyber incidents, the sector of the target firm, year-fixed effects, firm characteristics, and news sources. We cannot identify any time pattern in the magnitude and statistical significance of abnormal returns, as found by Gordon, Loeb, and Zhou (2011) in an earlier period. Moreover, some specific years have positive (albeit not statistically significant) abnormal returns. In contrast, we identify a significant marginal effect on firms in the health sector, translating into an average loss of USD 1.2 billion per event and for cyber incidents involving data breaches (average loss of USD 1.9 billion). In particular, the interaction coefficient, which measures the combined effect of the firms belonging to the intersection of the two subsamples, reaches as low as -7.07% and is statistically significant at the 10% level. In contrast, we cannot find evidence that ransomware attacks significantly affect firms' value, nor that the financial sector is more sensitive than others, among other non-statistically significant results. We next test whether intrinsic firm characteristics, typical performance determinants in the financial and accounting literature, can explain the variations in abnormal returns. Without concluding, we test the firm size, age, book-to-market ratios, and price-to-earnings ratios together. In only one specification, the firm size and the price-to-earnings ratio weakly explain abnormal returns, with firm size (price-to-earnings ratio) being a positive (negative) determinant, but with statistical significance levels well above the 1% level, a weak overall explanatory power (7.5% at most), and a large share of variance captured by the intercept. Finally, we test whether the source of the news has an impact on the abnormal returns. We control whether Reuters, Twitter, or another source first releases the news. Once again, the explanatory power does not allow us to reject the null of a particular effect.

²“Aggregated” stands for all cyber incidents considered together, irrespective of their type or of the industry to which the firm they target belongs to.

This paper contributes to the literature by providing a conceptual reinterpretation of the stock-market effects of cyber incidents under modern inference standards. While earlier studies document negative abnormal returns around cyber events, and more recent work notes that statistical significance weakens under improved estimation, no study has jointly combined (i) a modern, NLP-based identification of cyber incidents from high-frequency news, (ii) fully bias-adjusted event-study inference accounting for cross-correlation and event-induced variance, and (iii) a systematic decomposition of heterogeneity across incident types and sectors. Our results show that the disappearance of aggregate statistical significance is not merely a technical artifact of more conservative standard errors, but rather reflects a highly skewed distribution of effects in which economically meaningful market reactions are concentrated in specific settings—most notably data breaches affecting healthcare firms. This perspective reconciles apparently conflicting findings in the literature and clarifies when cyber incidents meaningfully affect firm value.

The remainder of the paper proceeds as follows. Section 2 introduces the existing literature and develops related hypotheses. Section 3 presents the data and methods, Section 4 details the results, and Section 5 concludes.

2. Literature review

2.1. Event study methodology and critical results

2.1.1. Methodology

Event study methods are heavily employed in all branches of economics. The first modern form of the empirical test to assess the impact of one or several economic events on one or several data series dates back to Ball and Brown (1968), whose research focuses on the impact assessment of changes in accounting measures on stock prices.³ Fama, Fisher, Jensen, and Roll (1969) also help define the modern empirical approach to event studies, as they study the abnormal returns around stock split announcements to falsify the efficient market hypothesis. Although they find significant positive abnormal returns around these announcements, with a market model, they also rationalize them.

³Dolley (1933) is alternately cited as the first author to employ event studies on stock splits.

While studying the effects of the accuracy of 336 forecast annual earnings disclosures on stock prices over the 1963–1967 period, Patell (1976) introduces a first adjustment in the event study methods. He proceeds by scaling the standard errors used for the t-test of significance by the standard deviation of the estimation period residuals. Boehmer et al. (1991) propose a similar adjustment to scale standard errors based on the event-induced variance (see also, Savickas, 2003). The adjustments above are the first to account for the event-induced heteroskedasticity of the time series under consideration.

On the other hand, empirical adjustments may be needed to account for the cross-sectional dependence of several firms affected by the same event. To alleviate this effect, Malatesta (1986) proposes the joint generalized least squares approach. Using simulations, he cannot demonstrate, however, that this approach dominates more simplistic cross-sectional dependence treatment. Similarly, Salinger (1992) takes the standpoint that actual abnormal returns must be uncorrelated even though their estimates are not. He derives an adjustment formula for the standard errors using a standard market model approach and a dummy variable. Applying this correction to an event study of post-merger performance, he shows that omitting the adjustment can lead to severe over-rejection of the null hypothesis of no abnormal returns. Similarly, Karafiath (1994) develops the dummy variable approach, which consists of appending a vector of zeros (ones) outside (inside) the event window. This approach enables obtaining abnormal returns in a single step, allowing for more straightforward adjustments of the standard errors and an easier interpretation of the residuals.

Another drawback of event studies is that they often rely on daily returns, which are more prone to autocorrelation, leading to an underestimation of standard errors. Brown and Warner (1985) propose a series of adjustments to address this issue and discuss the cross-sectional dependence and event-induced variance issues.

The studies above and the state-of-the-art methods generally apply to abnormal returns of short windows around the event considered (CARs). However, another strand of literature recognizes issues with estimating longer windows around and after the events, known as buy-and-hold abnormal returns (BHAR). Advances in treating such BHAR are discussed in Campbell, Lo, and MacKinlay (1997) and Kothari and Warner (2007).

Finally, Kolari and Pynnonen (2010) develop a t-statistic that summarizes and treats all the aforementioned CAR adjustments proposed. It deals jointly with event-induced variance, cross-correlation, and autocorrelation biases. They test the unbiasedness of their standard

errors and the power of the tests in simulations and generate a test that is the most effective available parametric test to date.⁴⁵ An additional justification for cross-correlation and induced variance correction stems from empirical evidence of spillovers surrounding data breaches and cyber incidents at various levels and across industries. For instance, Eckert, Gatzert, and Schubert (2023) uncover a spillover effect in the cyber insurance industry around data breaches, and Islam, Wang, Farah, and Stafford (2022) show that trading volume and uncertainties increase for rival control firms (peers) to those that cyberattacks have targeted.

2.2. *Cyber incident costs*

2.2.1. *Abnormal returns around cyber incidents*

Gordon et al. (2011) assess the impact of information security breaches on stock returns by estimating CARs over a three-day event window centered on news of cybersecurity incidents. They find that news about information security breaches generates significant negative CARs for publicly traded firms. They additionally uncover a significant downward shift in the impact of security breaches in the post-9/11 sub-period. They interpret these findings with a shift in investors' attitudes towards cyber breaches. In a similar study, Campbell et al. (2003) uncover highly significant (non-significant) negative CARs for information security breaches (not) involving unauthorized access to confidential data.

The studies above estimating the effects of cyber incidents on CARs do not adjust their results for the presence of cross-correlation (SUR estimations of Zellner, 1962) and do not adjust the standard errors for event-induced variance (Boehmer et al., 1991) and cross-correlation (Kolari and Pynnonen, 2010). They also find support for time-varying effects that we can test in a more recent subsample. Thus, we formulate our first null hypotheses as follows:

- H_{1a} : Cyber incidents do not generate economically significant CARs either with OLS estimation or with adjustments of event-study dummy coefficients for cross-correlation (SUR estimation).

⁴This literature review voluntarily skips the more minor literature on non-parametric tests. See, *e.g.*, the non-parametric rank test of Kolari and Pynnonen (2011).

⁵Also see Lee and Varela (1997) for tests of the different available standard errors adjustments.

- H_2 : The CARs' statistical significance is insensitive to adjustment that considers cross-correlation and event-induced variance.
- H_3 : Abnormal returns around cyber incidents are time-invariant.

Johnson, Kang, and Lawson (2017) also show that, on average, publicly traded firms in the U.S. face CARs of about -0.37% over a data breach. Additionally, breaches resulting from payment card fraud contribute more to negative abnormal returns than other breach types. Lastly, they find a positive link between the magnitude of these CARs and the size of the card breach. Lending, Minnick, and Schorno (2018) relate corporate governance and social responsibility to the probability of data breaches. Measuring changes in stock returns after breaches, they uncover a persistent effect of -3.5% of one-year BHAR. They also find that banks with breaches experience significant decreases in deposits, while non-banks experience large sales decreases. Tosun (2021) studies how financial markets react to unexpected corporate security breaches in both the short- and the long-run. He uncovers that market participants anticipate negative stock price changes while analyzing selling pressure and liquidity measures. However, the negative effect is significant only on the day following the public announcement of the security breach and is linked to an adverse impact on the firm's reputation. Kamiya, Kang, Kim, Milidonis, and Stulz (2021) find additional support for a reputation loss of targeted firms through a drop in credit ratings or a decrease in sales growth. Based on these results, we argue that our novel and updated dataset, along with the most advanced methods, calls for testing the following null hypothesis and sub-hypotheses,

- H_4 : CARs around cyber incidents cannot be explained by:
 - H_{4a} : Type of cyber incidents (ransomware, data breach, security breach, etc.)
 - H_{4b} : Type of sector the target firm belongs to (technology, health, financial, etc.)
 - H_{4c} : Typical firm characteristics such as size, book-to-market ratio, price-earnings ratio, and firm age

Andreadis, Kalotychou, Louca, Lundblad, and Makridis (2023) investigate how information dissemination about cyberattacks through major news sources affects municipalities' access to finance, particularly in the municipal bond market. Using a differences-in-differences framework, they show that an increase in the number of cyberattacks covered by news articles at the county level and the corresponding number of cyberattack news articles significantly

adversely impact municipal bond yields. A 1% rise in the number of cyberattacks covered in news articles results in offering yields increasing by 3.7 to 5.9 basis points, depending on the level of cyberattack exposure. In our case, testing for the number of news articles as a determinant of CARs would imply the use of forward-looking information and thus bias our results. However, with our sample of the first available news, we can test for the impact of the news source. Thus, we define our last null hypothesis as follows:

- H_5 : The news source does not drive the impact of cyber incidents on CARs.

Finally, other determinants of abnormal returns around cyber incidents have been used, for which our dataset does not allow testing. They include Jensen and Paine (2023), who use data on municipal IT investments, ransomware attacks, and bond performance. They cannot find an immediate impact on bond yields within 30 days of a cyberattack. However, over the subsequent 24 months following the cyberattack, municipal bond yields gradually decreased, while IT spending increased. They argue that this decline in bond yields is driven by a reduction in cyber risk due to increased IT investment. Gordon, Loeb, and Sohail (2010) examine how voluntary information security disclosures impact firm value, using a dataset of 1,641 firms that disclose such information and 19,266 that do not. They find that these disclosures can reduce litigation costs and lower a firm's cost of capital by reducing information asymmetry between management and investors. They find a positive relationship between voluntary information security disclosures and firm value. Firms that disclose this information also exhibit narrower bid-ask spreads than those that do not. Hilary, Segal, and Zhang (2016) also study cyber risk disclosures and uncover that the market reaction to cyber breaches is statistically significant but economically narrow. For an exhaustive literature review and meta-analysis of the impact of cyber-incidents on the stock market, also refer to Spanos and Angelis (2016).

2.2.2. Other methods for cyber incident costs estimation

Bouveret (2018) studies the global cyber risk for the financial sector and the various types of cyber incidents. Using a Value at Risk (VaR) framework, he uncovers an average country loss from cyberattacks of USD 97 billion and a VaR between USD 147 and 201 billion. He argues that essential potential aggregated losses exist in the financial sector, several orders of magnitude higher than the cyber insurance market could cover. Romanosky (2016) delves

into the nature and expenses related to cyber events. His analysis of a dataset comprising over 12,000 incidents reveals a skewed cost distribution, with an average cost of USD 6 million and a median cost of USD 170,000, akin to a typical annual IT security budget for firms. This leads to the hypothesis that firms may be operating optimally secure due to the relatively low costs, thus investing modestly in data protection measures. Other related studies include *e.g.*, Anderson, Barton, Böhme, Clayton, Eeten *van*, Levi, Moore, and Savage (2013); Hugues, Bohl, Irfan, Margolese-Malin, and Solorzano (2017); Anderson, Barton, Böhme, Clayton, Ganan, Grasso, Levi, Moore, and Vasek (2019).

Another strand of research in finance and cybersecurity uses asset pricing techniques and the cross-sectional analysis of listed stocks. Even though this literature relates more to the perceived risk than actual realization, this literature is worth mentioning. It includes *e.g.*, Florackis, Louca, Michaely, and Weber (2023); Jamilov, Rey, and Tahoun (2023); Jiang, Khanna, Yang, and Zhou (2023); Liu, Marsh, and Xiao (2022) who all find that the cyber risk is priced to some extent in the cross-section of stock returns.

A last and emerging strand of cyber incidents monitoring and cost estimation literature attempts to assess whether direct forecasts of future incidents are achievable through machine learning approaches. In particular, Almahmoud, Yoo, Damiani, Choo, and Yeun (2025) show that proactive machine learning and Bayesian GNN enable bias-free, data-driven prediction of cyber-attack trends and technology gaps, optimizing cybersecurity planning and resource allocation. Similarly, Celeny and Maréchal (2023) show that an increase of 10% in their NLP-derived cyber score commands a 10% increase in a firm’s likelihood to experience a cyber-incident over the next 12 months.

3. Data and methodology

3.1. Market data

We download public equity data from Wharton Research Data Services (WRDS), using the Center for Research in Security Prices (CRSP) and S&P Global Market Intelligence’s Compustat database. We report the list of variables in Table A1. We write a Python script that queries all available information from WRDS’ API and filters the firms so that all of the retained firms have at least one cyber incident relating to them in our news dataset

described below. We extract daily stock returns and financial ratios for 119 firms between January 2012 and December 2022. We also download the one-month Treasury bill rate and returns on the market, book-to-market (HML), and size (SMB) factors from the Kenneth French data repository⁶.

We exclude firms with market capitalization below USD 300 million to limit noise arising from thin trading, wide bid–ask spreads, and elevated volatility, which can distort abnormal-return inference in daily event studies. As documented in the event-study literature, micro-cap stocks often exhibit inflated variance and reduced reliability of standard errors. Including such firms would primarily increase estimation noise and make statistical significance less likely, implying that our main conclusions, particularly the lack of aggregate significance, would be reinforced rather than weakened.

3.2. *News data*

3.2.1. *News treatment*

We collect news headlines from the Refinitiv Eikon platform, which is recognized for its comprehensive financial news coverage and use in previous research. This platform categorizes news with a “Cybercrime” tag, focusing on cyber incidents. The taxonomy provided by Refinitiv for each incident type within this tag is more granular and will be detailed later (see Fig. 6). We begin with the download of 106,248 headlines. We first filter this dataset by keeping only English headlines related to listed companies, reducing our sample to 27,244. Second, we restrict the sample to North American firms using their CUSIPs and narrow it to 12,561. To identify characteristic keywords in cyber incident headlines, we compare 12,297 of these headlines against a standard news corpus using a volcano plot, as illustrated in Figure 1.

⁶Available at: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html




Fig. 1: **Volcano plot of cyber incident headlines against normal headlines**

Volcano plot of a normal news corpus against a cybercrime news corpus. The population of both samples is 12,297 and is extracted from Refinitiv. The x-axis represents a log2 fold change, while the y-axis depicts the negative logarithm (base 10) of the p-value derived from a chi-squared test. The x and y thresholds are set to select the relevant keywords optimally.

The volcano plot contrasts word frequencies between two distinct datasets. The x-axis shows the log2 fold change, indicating the relative frequency shift of words between cybercrime and regular headlines. Words more common in cybercrime headlines have a higher positive value. The y-axis presents the negative log (base 10) of the p-value from a chi-squared test, highlighting the statistical significance of the frequency differences. Words at the top of the plot show substantial frequency changes, thus distinguishing cybercrime headlines from regular news.

Based on this analysis, we identify keywords typically associated with cyber incidents, such as *breach*, *hacker*, *cybersecurity*, *ransomware*, *malware*, *leak*, *vulnerability*, *data*, and *attack*. These terms, though not exhaustive, cover a broad range of cyber incident scenarios (e.g., *[Company Name] is experiencing a [Keyword]*). Applying these keywords, we filter the headlines to include specific cyber terms and company names, making exceptions for Meta and Alphabet to capture incidents related to Google or Facebook. This process yields 3,606 headlines.

We then use the MoritzLaurer/DeBERTa-v3-large-mnli-fever-anli-ling-wanli model from the Hugging Face platform for text entailment.⁷ This is to validate headlines whose content is more complex. For instance, *[Company Name] has been under a cyber incident or attack*. We seek an entailment score above 93% to underfit the results for later manual evaluation. Based on BERT and DeBERTa, this model is fine-tuned on various natural language inference datasets, offering 885,242 hypothesis-premise pairs, and is highly rated on Hugging Face as of June 6, 2022. This step results in 1,465 headlines.

To construct a unique set of cyber-incident disclosures, we retain at most one headline per firm-day and require that consecutive incidents for the same firm be at least two months apart, with exceptions for corporations such as Meta, Alphabet, Google, and Apple due to their higher frequency of incidents. This procedure results in 368 headlines. We then manually verify these headlines for date accuracy, relevance, and incident type, yielding a final sample of 270 headlines, for which we record the target company identifiers, incident type, event date, and news source. We depict all filtering steps in Figure 2. Some cyber incidents generate multiple news headlines within the same trading day, particularly for large or high-visibility firms. To avoid double-counting the same information event, we treat the earliest verified headline per firm-day as the disclosure date. Additional same-day headlines are recorded but not treated as separate events, as they typically reflect the dissemination of already public information rather than new disclosures. This approach aligns the event date with the first public release of information and follows standard practice in news-based event studies.

⁷Available at <https://huggingface.co/MoritzLaurer/DeBERTa-v3-large-mnli-fever-anli-ling-wanli>




Fig. 2: **Filtering process of cyber incident headlines**

This diagram depicts the filtering steps of the original cyber incident headline set from Refinitiv. We only keep headlines in English that mention listed companies with an identifier (RIC). Next, we keep North American listed companies that have a CUSIP. We filter based on cyber incident keywords. We subsequently check whether the headline is a premise to the following entailment question: “[Company Name] has been under a cyber incident or a cyberattack.” We additionally keep the earliest headline per day and company. These must be two months apart for a given company (starting with the earliest). Finally, we manually change the dates and drop headlines that are not about a cyber incident or are not the first to mention it.

We further filter the number of events by merging them with the WRDS and Compustat databases. Hence, we drop all firm-related events not in the WRDS database. We also drop firms not listed on the event day, and firms where the event happens right before or after a long weekend.⁸ We end up with 167 events relating to 73 firms. These filtering steps are illustrated in Figures 3 and 4.

⁸We follow standard recommendations in event study as we drop firms with small market capitalization, see, *e.g.* Dimson and Marsh (1986).




Fig. 3: **Evolution of the number of incidents**

In step 1, we drop all firms not in the WRDS database. In step 2, we drop all incidents affecting cybersecurity-providing firms. In step 3, we drop incidents where the firm’s market capitalization is lower than \$300 million, the firm is not listed during the incident, the incident is in 2023, or the incident window is larger than five days. In step 4, we drop all firms not listed between 2013.12.19 and 2022.10.13 and/or whose cyber incident(s) did not occur between those dates.




Fig. 4: **Evolution of the number of firms**

In step 1, we drop all firms not in the WRDS database. In step 2, we drop all incidents affecting cybersecurity-providing firms. In step 3, we drop incidents where the firm’s market capitalization is lower than \$300 million, the firm is not listed during the incident, the incident is in 2023, or the incident window is larger than five days. In step 4, we drop all firms not listed between 2013.12.19 and 2022.10.13 and/or whose cyber incident(s) did not occur between those dates.

We depict the sector and attack type distributions in Figures 5 and 6, respectively. Most firms are from the technology sector, whereas the types of attacks are more diversified. The dataset’s two most common types of attacks are data breaches and software breaches.




Fig. 5: **Sector distribution**

Sector distribution of the firms remaining after step 3 on Figure 4. We use Refinitiv’s sector classification.




Fig. 6: **Attack distribution**

Distribution of the different types of attacks based on the Refinitiv classification of incidents remaining after step 3 on Figure 3. The “Other” category regroups the Cyber breach (1.93%) and Stolen funds (1.93%) categories.

3.2.2. Abnormal returns

We compute abnormal returns with the Fama and French (1992) three-factor model (FF3) using Eq. 1 below.

$$R_i^e = \alpha_i + \beta_i^{Mkt} R_{Mkt}^e + \beta_i^{SMB} R_{SMB} + \beta_i^{HML} R_{HML} + \sum_{j=0}^k AR_{event_j} \mathbb{1}_{event_j} \quad (1)$$

where R_{Mkt}^e is the excess return on the market portfolio, R_{SMB} is the return on the size factor, R_{HML} is the return on the book-to-market factor, k is the total number of events in the dataset, AR_{event_j} is the abnormal return on event j and $\mathbb{1}_{event_j}$ is a dummy vector that takes the value “1” on the day of event j and “0” otherwise. This equation can be written as $R_i^e = X\beta_i$ in matrix notation. We estimate Eq. 1 for each firm, using Ordinary Least Squares (OLS) regressions and Seemingly Unrelated Regressions (SUR) (see Zellner, 1962). In the SUR estimation, the error terms are assumed to be correlated across the equations, and each equation must have the same number of observations. Hence, we perform SUR on a subset of the data that includes the most cyber incidents related to firms with daily returns for all trading days in the subset. We find the optimal subset to start on December 19, 2013, and end on October 13, 2022. There are 126 cyber incidents relating to 48 firms in this subset of the data, as can be seen in Figures 3 and 4. We also compute abnormal returns with the zero-benchmark model presented in Eq. 2.

$$R_i^e = \alpha_i + \sum_{j=0}^k AR_{event_j} \mathbb{1}_{event_j} \quad (2)$$

We compute CARs over a three-day window centered on cyber incidents.

3.2.3. *t*-stat adjustments

We use scaled abnormal returns and two t-test statistics from Kolari and Pynnonen (2010). We borrow their notations for the remainder of the paper. The scaled abnormal returns can be computed using Eq. 3,

$$A_{i,t} = \frac{AR_{i,t}}{s_i \sqrt{1 + x'_t(X'X)^{-1}x_t}} \quad (3)$$

where s_i is the regression residual standard deviation, X is the matrix of explanatory variables from the matrix notation of Eq. 1, and x_t is the t -th row of X . A feasible estimator of the variance of the scaled abnormal returns can be computed using Eq. 4,

$$s_A^2 = \frac{s^2}{1 - \bar{r}}, \quad (4)$$

where s^2 is the sample cross-sectional variance of event-day scaled abnormal returns and $\bar{r}$ is the average sample cross-correlation of the residuals.

The ADJ-BMP statistic is computed as shown in Eq. 5, below.

$$t_{AB} = \frac{\bar{A}\sqrt{n}}{s_A \sqrt{1 + (n-1)\bar{r}}}, \quad (5)$$

where n is the number of firms in the sample. This statistic is residual cross-correlation and event-induced volatility-adjusted.

The ADJ-PATELL statistic is computed as shown in Eq. 6.

$$t_{AP} = \frac{\bar{A}\sqrt{n}}{\sqrt{(m-p-1)/(m-p-3)} \sqrt{1 + (n-1)\bar{r}}}, \quad (6)$$

where p is the number of explanatory variables in the expected return regression (three in our case) and m is the number of days in the sample. This statistic is residual cross-correlation-adjusted.

Event studies based on daily returns may be affected by serial correlation. However, this concern is mitigated in our setting for several reasons. First, our analysis focuses on short event windows, for which prior work shows that autocorrelation effects are limited (Brown and Warner, 1985). Second, expected-return models are estimated over long non-event peri-

ods, and the Fama–French factors absorb much of the predictable variation in returns. Third, the Kolari and Pynnonen (2010) adjustments explicitly address cross-sectional dependence and event-induced variance, which are the dominant sources of biased inference in multi-event event studies. Taken together, these features ensure that residual serial correlation does not materially affect our statistical conclusions.

4. Results

4.1. *CAR estimations*

Figure 7 presents the correlation matrix of the residuals obtained from estimating Eq. 1. Overall, the residuals do not have high correlations; the average correlation is 0.012. Since the SUR estimation is equivalent to OLS when the error terms are uncorrelated between equations, we expect the results obtained with the SUR estimation to be similar to the ones obtained with OLS.




Fig. 7: **Correlation of residuals**

Correlation matrix of residuals, obtained by estimating Eq. 1.

Table 1 presents the average CARs and unadjusted and adjusted t-statistics.

	Zero-benchmark			FF3 benchmark		
	OLS	OLS limited	SUR	OLS	OLS limited	SUR
$\overline{CAR}$	-0.89%	-0.88%	-0.79%	-0.74%	-0.84%	-0.78%
Unadjusted t-stat	-1.949	-2.258	-2.325	-1.856	-2.579	-2.434
ADJ-PATELL	-	-0.704	-0.455	-	-2.095	-1.807
ADJ-BMP	-0.157	-0.323	-0.273	-0.698	-1.226	-1.033
Number of incidents	167	126	126	167	126	126
Number of days in regression	-	2219	2219	-	2219	2219

Table 1: **Average CAR and t-stats**

$\overline{CAR}$ is the average cumulative abnormal return over a three-day window around the incidents, relative to the zero benchmark model (shown in Eq. 2) or the three-factor model of Fama and French (1992) (shown in Eq. 1). ADJ-BMP and ADJ-PATELL refer to the adjusted statistics of the same name from Kolari and Pynnonen (2010). The ADJ-PATELL statistic is cross-correlation-adjusted, and the ADJ-BMP statistic is cross-correlation and volatility-adjusted. The OLS model allows for the number of days in the regression to differ for each firm, and the average number of days is 2398. The OLS limited model is the OLS model restricted to the same period, the same subset of firms, and the same incidents as the SUR model.

We observe that the average CARs obtained with the SUR approach are close to the ones obtained with OLS. The average CAR for the three-factor benchmark is economically significant, at around -0.8%, and statistically significant, at least at the 10% level, when using unadjusted t-statistics. These results are close to the ones obtained in the existing literature over the past twenty years, and in particular, to the most recent event study on cyberattacks by Kamiya et al. (2021), who find an average CAR of -0.768% for the three-factor benchmark. Thus, we reject our null hypothesis H_1 regarding the economic significance of CARs. The t-statistics decrease when adjusting for residual cross-correlation and event-induced volatility. Importantly, none of the average CARs are statistically significant using the ADJ-BMP statistic, that is, once all biases are accounted for. Hence, the previously claimed statistical significance of CARs around cyberattacks seems not to hold once we properly adjust standard errors for the induced market-level effects of the cyber incidents. We are therefore not able to reject our hypothesis H_2 .

We further investigate the evolution of abnormal returns around cyber incidents. Figure 8

shows the average abnormal return for each day in an 11-day window around cyber incidents. We observe that the abnormal returns drop up to three days after the day of the cyber incident but recover soon after. However, this is not a price recovery since post-event returns remain around zero. Figure 9 shows that the CARs are at zero at the end of the window, but this is more due to a systematic increase before the event rather than a price recovery.⁹

Thus, it is difficult to conclude whether cyber incidents are informative regarding the firm’s prospects or if they are purely temporary for technical (or behavioral) reasons, as advocated by Shleifer and Vishny (1997). While this is not the primary focus of the study, it further supports our claim that the significance of cyber incidents may not have as dramatic an effect on the firm as previously found in some research.

The presence of mild positive abnormal returns before the event date is unlikely to reflect information leakage or systematic misalignment of event timing. Event dates are defined using the earliest verified public disclosure, and clustered firm-level events are excluded by construction. Moreover, the persistence of this pattern under both the three-factor benchmark and the zero benchmark suggests that it does not arise from benchmark misspecification. A more plausible explanation relates to sample composition: cyber incidents disproportionately affect firms in high-growth sectors, particularly technology, which experienced sustained positive momentum over much of the sample period. Prior work on corporate disclosure timing also suggests that adverse information may be released during periods of strong performance, further contributing to this pattern. Importantly, these pre-event dynamics do not affect the interpretation of the estimated event-window abnormal returns.

⁹To verify that benchmark misspecification does not drive the observed pre-event abnormal returns, we estimate abnormal returns using several alternative benchmarks, including a zero-return benchmark. The pre-event pattern is robust across these specifications, indicating that it is not an artifact of the expected-return model. We therefore interpret this pattern as reflecting sample composition rather than information leakage or event-date misalignment.




Fig. 8: **Average abnormal returns**

Average abnormal daily returns in an 11-day window centered on the cyber incidents. Day 0 is the day of the cyber incident. The OLS limited model is the OLS model restricted to the same period, the same subset of firms and incidents as the SUR model. Abnormal returns are computed with the three-factor model of Fama and French (1992), following Eq. 1.




Fig. 9: **Cumulative average abnormal returns**

Cumulative average abnormal daily returns in an 11-day window centered on the cyber incidents. Day 0 is the day of the cyber incident. The OLS limited model is the OLS model restricted to the same period, the same subset of firms and incidents as the SUR model. Abnormal returns are computed with the three-factor model of Fama and French (1992), following Eq. 1.

4.2. CARs through time

To test our hypothesis H_3 regarding the time invariance of CARs around cyber incidents, we add year fixed-effect dummies in a panel regression of CARs. We use the CARs arising from the OLS and SUR estimations to ensure our results would not be driven by potential residual cross-correlation. We report our results in Table 2.

Dependent variable: CAR		
	OLS	SUR
	Model 3	Model 3
2013	0.557 [0.336]	- -
2014	-0.029 [-0.017]	- -
2015	-1.109 [-0.598]	-0.899 [-0.659]
2016	-1.019 [-0.727]	-1.760* [-1.756]
2017	-0.324 [-0.262]	-0.662 [-0.709]
2018	-1.536 [-1.243]	-1.790** [-1.981]
2019	-0.651 [-0.430]	-0.028 [-0.022]
2020	-0.268 [-0.260]	0.446 [0.550]
2021	-1.799 [-1.456]	-1.527 [-1.581]
2022	-1.716* [-1.822]	-0.228 [-0.309]
Observations	167	126
R-squared	0.050	0.084

Table 2: **Determinants of CARs - Incident year**

Results of regressions of the CARs on years. T-statistics are reported in brackets. The CARs are multiplied by 100. The first column uses the CARs obtained using OLS regressions, and the other column uses the CARs obtained with SUR. *, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

We cannot identify any monotonic increase or decrease of CARs magnitude through time, as previously advocated by *e.g.* Gordon et al. (2011), who find a structural decrease in the effect in their sample. Moreover, in the OLS specification (with 167 observations), we find that the coefficient for 2013 is positive, whereas, for the SUR estimation, that of 2020 is positive, albeit not statistically significant. Moreover, only two yearly dummy coefficients are statistically significant at the 10% level (2016 for SUR and 2022 for OLS), and only

one is significant at the 5% level (2018 for SUR). Finally, the resulting R^2 is small in both estimations, with 5% for the OLS case and 8.4% for the SUR. These additional pieces of evidence make us confident that no time structure is present in the effects of cyber incidents on firms' valuation. Thus, we are not able to reject our null hypothesis H_3 .

4.3. Event type and target sector

To test the hypotheses H_{4a} and H_{4b} , we adopt the same panel regression setting and regress CARs on the type of cyber incidents reported in the news headlines, the sector the firms belong to, and their interactions. We report our results in Table 3.

Dependent variable: CAR				
	OLS		SUR	
	Model 1	Model 2	Model 1	Model 2
Data breach	-1.287* [-1.907]	0.261 [0.167]	-1.178** [-2.106]	-1.220 [-1.013]
Software breach	-0.380 [-0.480]		-0.178 [-0.290]	
Cyber breach	1.152 [0.438]		1.270 [0.349]	
Social breach	0.081 [0.045]		-0.086 [-0.061]	
Ransomware	0.395 [0.258]		-0.851 [-0.705]	
Shutdown	-0.436 [-0.041]		-0.753 [-0.747]	
Twitter breach	-0.846 [-0.070]		-0.683 [-0.817]	
Facebook breach	2.378 [1.010]		1.786 [1.096]	
Stolen funds	0.362 [0.135]		-0.617 [-0.231]	
Mitigation	1.762 [0.748]		1.764 [1.083]	
Other		1.362 [0.960]		0.422 [0.360]
Technology		-2.172 [-1.413]		-0.923 [-0.732]
Consumer products		1.561 [0.793]		-0.272 [-0.166]
Financials		-3.825 [-0.722]		0.955 [0.856]
Healthcare		-5.210 [-1.594]		1.399 [0.500]
Industrials		-1.713 [-0.680]		-3.471 [-1.455]
Data breach * Technology		0.434 [0.285]		0.329 [0.270]
Data breach * Consumer products		-2.341 [-1.173]		0.599 [0.380]
Data breach * Financials		3.799 [0.799]		0.955 [0.856]
Data breach * Healthcare		-0.251 [-0.080]		-7.069* [-1.917]
Data breach * Industrials		-1.381 [-0.524]		3.966* [1.708]
Observations	167	167	126	126
R-squared	0.042	0.106	0.085	0.109

Table 3: **Determinants of CARs - Incident type and target sector**

Results of regressions of the CARs on the firm's sector and the type of the incident. t-statistics are reported in brackets. The CARs are multiplied by 100. "Consumer products" regroups the consumer cyclical and consumer non-cyclical sectors. The first two columns use the CARs obtained using OLS regressions, and the other two columns use the CARs obtained with SUR. *, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

In Model 1, we report the results of CARs regressed in a panel on dummy vectors set to “1” (“0”) when the incident belongs (does not belong) to the type of cyber incidents under scrutiny. The most significant incident affecting firm returns is the data breach, with an economic magnitude of -1.287% (-1.178%) with CARs estimated with the OLS (SUR) approach. Over the full sample, it translates into total, average, and median losses of USD 129 billion, 1.9 billion, and 105 million, respectively. The data breach coefficient is also the single one to be statistically significant at the 10% level (5% level) for CARs estimated with OLS (SUR). Not only do none of the other events seem to explain the CARs statistically, but they also show interesting coefficient sizes. For instance, the “cyber breach”, “social breach”, “ransomware”, “Facebook breach”, “stolen funds”, “mitigation”, as well as the “other” categories, all show positive signs for their coefficient. This coefficient magnitude reaches even 2.378% for the “Facebook breach” type of incident (*i.e.* when a company’s Facebook account has been compromised) for CARs estimated with OLS. In Model 2, we use the same approach and regress CARs in a panel on dummy vectors coded “1” when the firm belongs to a specific sector (technology, consumer products, financials, healthcare, and industrials) as well as their interaction with the data breach dummy vector. Given the few remaining degrees of freedom, we restrict the interaction terms with this type of cyber incident because the data breach coefficient is the only statistically significant in the Model 1 specification. We find that none of the sector coefficients is statistically significant, with the healthcare sector coefficient being the most negative at -5.21% and approaching the 10% statistical significance level for the CARs arising from the OLS estimation.¹⁰ Over the full sample, it translates into total, average, and median losses of USD 12 billion, 1.2 billion, and 65 million, respectively.

The financials sector ranks second with -3.83%, followed by technology (-2.17%) and industrials (-1.71%). Conversely, the consumer products sector coefficient is positive at 1.56%. Model 2, estimated with CARs estimated with the SUR approach, shows a different pattern. First, the sign of the coefficients of three sectors switches. In particular, that of the healthcare sector becomes positive. However, this is largely because the effect is captured by the interaction with the data breach dummy vector, which reaches -7.07% and

¹⁰In an unreported test, we also regress CARs only on the sectors, dropping the interaction terms. The coefficients are virtually the same, except for the healthcare sector coefficient, which reaches the usual statistical levels. These results are available upon request.

is statistically significant at the 10% level. This supports the view that the marginal effect of having a data breach for a healthcare company is the most detrimental situation a company can bear around a cyber incident. The information content does not allow us to conclude why this specific combination is so damaging. However, we can hypothesize that investors' views about potentially compromised private clinical data are particularly affected. This specification also yields the highest R^2 close to 11%. Given the absolute contribution of the data breach effect versus the other types of incidents and its relative impact when combined with the healthcare sector, we find support to reject our null hypotheses H_{4a} and H_{4b} . These results also align with those of Gordon et al. (2011), who identify that data breaches with availability concerns are the most prone to generate negative abnormal returns.

Although our dataset does not include a continuous numerical severity score for cyber incidents, the incident-type taxonomy provided by Refinitiv serves as a meaningful proxy for the magnitude of events. Categories such as data breaches, shutdowns, or stolen-funds incidents typically reflect substantially more severe disruptions than social-media account compromises or credential-related incidents. The strong and statistically significant effects we document for data breaches, therefore, indicate that incident type already captures an important severity dimension. Interpreting heterogeneity through this lens reinforces our conclusion that only the most severe and information-sensitive cyber incidents generate economically meaningful market reactions.

4.4. *Firm characteristics*

We now study the firm characteristics as explanatory variables to test our hypothesis H_{4c} . We repeat the panel data regression using continuous variables for those characteristics. We report the results in Table 4.¹¹

¹¹We only report the OLS for the full sample (150 observations) and the SUR estimation (117 observations) for the restricted sample. The small statistical significance observed for the constant and price-to-earnings ratio only comes from the observation restrictions. In unreported OLS tests on the restricted sample, we also observe this significance arising. These results are available upon request.

Dependent variable: CAR		
	OLS	SUR
	Model 5	Model 5
Constant	-6.040	-7.872*
	[-1.177]	[-1.682]
Firm size (ln)	0.206	0.348*
	[0.994]	[1.881]
Firm Age (ln)	0.155	-0.414
	[0.741]	[-0.946]
Book to market	-0.065	-0.035
	[-0.644]	[-0.247]
Price to earnings	0.001	-0.015**
	[0.115]	[-2.164]
Observations	150	117
R-squared	0.017	0.075

Table 4: **Determinants of CARs - Firm characteristics**

Results of regressions of the CARs on firm characteristics. t-statistics are reported in brackets. The CARs are multiplied by 100. The first column uses the CARs obtained using OLS regressions, and the other column uses the CARs obtained with SUR. Certain observations were dropped due to missing accounting data. *, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

The explanatory power of the considered characteristics is limited, with 1.7% for the CARs estimated with OLS and 7.5% for those estimated with the SUR approach. Only two variables of interest are statistically significant at the 5% level (price-to-earnings ratio) and the 10% level (firm size) in the SUR specification. To summarize these results, the larger the market capitalization of a firm but the smaller its price-to-earnings ratio, the more resilient it would be to cyber incidents. Nonetheless, most variation is unexplained and captured by the constant at -6.04% and -7.87% in the OLS and SUR specifications, respectively.

4.5. *News source effect*

To test our hypothesis H_5 , we regress CARs on the type of news source to test whether this influences the economic magnitude and statistical significance of the costs of cyber incidents. We report our results in Table 5.

Dependent variable: CAR		
	OLS	SUR
	Model 4	Model 4
Reuters	0.186 [0.219]	-0.477 [-0.718]
Twitter	-0.843 [-0.313]	-1.029 [-0.549]
Other	-0.840 [-1.239]	-0.472 [-0.927]
Observations	167	126
R-squared	0.001	0.006

Table 5: **Determinants of CARs - News source**

Results of regressions of the CARs on the news source. t-statistics are reported in brackets. The CARs are multiplied by 100. The first column uses the CARs obtained using OLS regressions, and the other column uses the CARs obtained with SUR.

We have 270 news on the unrestricted sample, from which we can estimate CARs with OLS, which is split between 156 news provided by Reuters, 15 by Twitter, and the rest classified as “other”. Different source types are extremely diverse and generally count as one. The split for the restricted sample of 126 CARs is 72 for Reuters and 4 for Twitter, the remainder classified as “other”. We find almost no explanatory power with these panel regressions, with an R^2 at 0.1% for CARs arising from OLS and 0.6% for those deriving from the SUR estimation. Moreover, we cannot find any significant coefficients at the usual statistical levels for the type of news source. Interestingly, the coefficient for Reuters is positive when CARs from OLS are regressed. In contrast, the most considerable negative magnitude is found for news from Twitter for CARs from both estimations, with -0.84% and -1.03% for CARs from OLS and SUR estimations, respectively. Thus, we cannot reject our null hypothesis H_5 , which supports the view that the source of the first available news does not affect CARs’ magnitude.

4.6. Robustness tests

We repeat the core analysis with a zero-benchmark model to ensure a misspecified benchmark does not drive our results. We present these results in Table 1, Figure A1, and Figure

A2. The results and significance are very similar to the ones obtained with the three-factor benchmark, which makes us confident that our benchmarks do not generate bias. An alternative way to consider these results is that the procedure generates the actual “raw” return that an average agent would experience while being invested in these stocks over the cyber incident windows.

5. Discussion

5.1. *Hypotheses explaining our results*

Several factors could explain our findings regarding the health sector and the data breach. However, these explanations are hypothetical at this stage and can only be based on the patterns observed in our results. First, data breaches involve unauthorized access to sensitive information, often leading to compromised personal or confidential data. The health sector deals with highly sensitive patient information, making data breaches particularly damaging. The potential for compromised clinical data, patient records, or personal health information may severely affect individuals and healthcare organizations. Second, the health sector is subject to strong regulations and compliance standards, such as the US’s 1996 Health Insurance Portability and Accountability Act (HIPAA).¹² The regulatory environment demands a high level of data protection. Any failure to safeguard patient information can result in significant legal and financial consequences; see, *e.g.*, Khansa, Cook, James, and Bruyaka (2012). This regulatory scrutiny amplifies the impact of data breaches in the health sector compared to other industries, which may also explain why the combined effect of data breaches on the health sector’s firms is even more detrimental. Third, health data is often considered more valuable than other information on the dark web due to its potential for various malicious activities, including identity theft and healthcare fraud. As a result, cybercriminals may specifically target the health sector to gain access to valuable data, contributing to the sector’s increased vulnerability. Finally, data breaches involving sensitive patient information can significantly erode public trust, leading to a permanent price depreciation (negative CARs). The potential harm to an organization’s reputation and the loss of patient confidence can have lasting effects, influencing investors’ perceptions and

¹²See, <https://www.govinfo.gov/content/pkg/CRPT-104hrpt736/pdf/CRPT-104hrpt736.pdf>

contributing to a more pronounced negative impact on market valuation. In contrast, the non-significant CARs we observe on ransomware attacks may be explained by their impact being more immediate and operational rather than directly affecting a firm’s valuation. Organizations with robust backup and recovery mechanisms can mitigate the financial losses associated with such attacks. Additionally, the financial demands of ransomware may not necessarily translate into a direct and lasting impact on a firm’s market valuation, explaining the observed lesser significance.

It is important to note that these explanations are speculative. Further research would be needed to validate and refine these hypotheses, considering additional factors and potential interactions that may influence the impact of cyber incidents on different types of firms and sectors.

5.2. *Behavioral vs. fundamental channels in market reactions to cyber incidents*

A central issue in interpreting market responses to cyber incidents is whether abnormal returns reflect fundamental valuation effects (*e.g.*, expected regulatory penalties, litigation risk, or operational disruptions) or behavioral mechanisms (*e.g.*, investor attention, salience, and short-term overreaction). Addressing this distinction is essential for understanding why aggregate effects weaken under rigorous inference while strong localized effects remain. The short-lived negative abnormal returns observed for most cyber incidents are consistent with a behavioral channel. Minor or moderately disruptive events often attract disproportionate media attention relative to their long-term cash-flow implications, generating temporary price pressure that dissipates as information is processed. Such patterns are well-documented in settings where salient news induces short-term overreaction followed by normalization. In contrast, the economically large and persistent negative effects associated with data breaches affecting healthcare firms are more plausibly driven by fundamental considerations. Healthcare organizations operate under strict regulatory regimes, face heightened legal liability following data breaches, and manage highly sensitive personal information. The compromise of such data can therefore generate lasting reputational damage and durable changes in firms’ risk profiles, which rational investors are likely to incorporate into prices. Taken together, our findings support a dual-channel interpretation of market reactions to cyber

incidents: behavioral responses dominate for less severe events, while fundamental valuation effects emerge in settings characterized by strong regulatory exposure and high information sensitivity.

5.3. Policy recommendations

Given the significant negative impact on firms in the health sector, policymakers should consider sector-specific cybersecurity guidelines and incentives to mitigate the vulnerabilities identified in the study. This could include tailored regulations and support for health-related organizations to enhance their cybersecurity posture.

Policymakers should encourage and enforce robust data protection measures through sector-specific regulations addressing the unique challenges of data breaches. They should also promote mechanisms for improved information sharing among firms within the same sector. We observe that firms from all sectors suffer from data breaches; therefore, encouraging collaboration and sharing best practices may enhance firms' overall cybersecurity resilience. Next, since we find that cyber incidents do not uniformly impact all firms, policymakers should collaborate with the insurance industry to develop contracts that specifically mitigate the impact of each incident. This could involve tailoring insurance coverage and premiums based on the sector, type of cyber incident, and other relevant factors identified in the study.

Last, increased collaboration between government agencies, private organizations, and cybersecurity experts should address the challenges of cyber incidents. The policy approach should be dynamic and targeted, considering the nuanced nature of cyber threats. Policymakers should address specific vulnerabilities identified in the study and collaborate with industry stakeholders to comprehensively respond to the evolving nature of cyber incidents.

5.4. Cybersecurity investment recommendations

Given the significant negative impact identified for firms in the health sector, there is a strong case for these organizations to contract cyber insurance products. Insurance policies should be tailored to cover potential economic losses associated with data breaches.

While our results are not statistically significant for firms in sectors other than healthcare, it is still advisable for all firms to consider cyber insurance. Coverage should be structured

to align with sector-specific risks and exposures. Standard cyber insurance policies may be sufficient for firms outside the health sector, whereas firms in the health sector should prioritize investments in cybersecurity measures aimed at preventing and mitigating data breaches. Such measures include advanced encryption technologies, access controls, employee cybersecurity training, and continuous monitoring to enhance threat detection and response. Overall, our results highlight the heterogeneity of cyber-incident impacts across sectors, suggesting that cybersecurity investments and insurance coverage should be tailored to the specific risk profiles faced by different industries.

Beyond managerial and policy implications, an important question is whether the abnormal returns associated with cyber incidents are economically exploitable. Several features of cyber-incident disclosures limit their practical tradability. News about cyber incidents diffuses rapidly across information channels, often within the same trading session, while bid-ask spreads tend to widen around such announcements, increasing execution costs. Moreover, for most incident types, abnormal returns reverse quickly, leaving little scope for systematic trading strategies. These considerations suggest that the documented return patterns are consistent with informationally efficient markets rather than persistent mispricing.

6. Conclusion

Our study delves into the impact of cyber incidents on listed firms, employing advanced event study methods and adjusted returns. By utilizing newswire headlines filtered for relevant information on cyber incidents, we use state-of-the-art estimation methods, specifically a seemingly unrelated regression estimation, to ensure unbiased cumulative abnormal returns coefficients. Our findings reveal an overall adverse effect of approximately -0.89%, regardless of the target firm or the type of cyber incident.

We adjust standard errors using corrections for event-induced variance and cross-correlation to address a crucial caveat in the existing literature. Unadjusted standard errors result in cumulative abnormal returns with significance below the 1% level, but our correction reveals no statistical significance at the aggregated level. This contradicts the prevailing consensus in event study literature. A more detailed analysis, incorporating variables such as the type of cyber incident and the sector of the target firm, reveals a significant marginal effect on

firms in the health sector and those involved in data breaches. In contrast, counterintuitively, ransomware or other attacks do not significantly affect firms' value. No heightened sensitivity is observed in the financial sector. Conversely, alternative explanatory variables such as time-fixed effects, firm characteristics, or news sources yield little to no explanatory power and statistical significance.

One alternative benchmark to be considered would be the “peer firms benchmark”, *i.e.* building the counterfactual returns using a set of firms with similar characteristics. This would help reduce the systematic pre-event positive returns we observe. However, we would be limited in our setting, as firms with characteristics identical to those subject to cyber incidents would also be more frequently affected. Another possible extension that would allow the research to be more conclusive on the permanent *transitory* price pressure debate is to study the effects of cyber incidents on returns and other market metrics such as turnover, volatility, liquidity (bid-ask spread), etc.

By adopting a standard error correction approach, utilizing an up-to-date dataset that reflects changes in cyber incident frequency, and employing a new set of headlines filtered with NLP methods, we argue that cyber incidents have less systematic detrimental effects on firms' value than previously claimed when considered in aggregate. However, we identify some edge cases, mainly when cyber incidents involve data breaches and when target firms belong to the healthcare sector, which is particularly detrimental to the firm. Our results also highlight the heterogeneity of the impact of cyber incidents, and not only call for improved investments and insurance contracts for the health sector regarding their data, but also for taking into account each specific risk across each sector.

References

- Almahmoud, Z., Yoo, P. D., Damiani, E., Choo, K.-K. R., Yeun, C. Y., 2025. Forecasting cyber threats and pertinent mitigation technologies. *Technological Forecasting and Social Change* 210, 123836.
- Anderson, R., Barton, C., Böhme, R., Clayton, R., Eeten *van*, M. J. G., Levi, M., Moore, T., Savage, S., 2013. Measuring the cost of cybercrime. *Workshop on the Economics of Information Security* 11, 265–300.
- Anderson, R., Barton, C., Böhme, R., Clayton, R., Ganan, C., Grasso, T., Levi, M., Moore, T., Vasek, M., 2019. Measuring the changing cost of cybercrime. *Workshop on the Economics of Information Security* 18, 1–32.
- Andreadis, L., Kalotychou, E., Louca, C., Lundblad, C. T., Makridis, C., 2023. Cyberattacks, media coverage and municipal finance. Available at <https://dx.doi.org/10.2139/ssrn.4473545>
- Ball, R., Brown, P., 1968. An empirical evaluation of accounting income numbers. *Journal of Accounting Research* 6, 159–178.
- Boehmer, E., Musumeci, J., Poulsen, A. B., 1991. Event-study methodology under conditions of event-induced variance. *Journal of Financial Economics* 30, 253–272.
- Bouveret, A., 2018. Cyber risk for the financial sector: A framework for quantitative assessment. Available at <http://dx.doi.org/10.2139/ssrn.3203026>
- Brown, S. J., Warner, J. B., 1985. Using daily stock returns: The case of event studies. *Journal of Financial Economics* 14, 3–31.
- Campbell, J. Y., Lo, A. W., MacKinlay, A. C., 1997. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, New Jersey.
- Campbell, K., Gordon, L. A., Loeb, M. P., Zhou, L., 2003. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Cybersecurity* 11, 431–448.

- Celeny, D., Maréchal, L., 2023. Cyber risk and the cross section of stock returns. Available at <http://dx.doi.org/10.2139/ssrn.4587993>
- Dimson, E., Marsh, P., 1986. Event study methodologies and the size effect: The case of UK press recommendations. *Journal of Financial Economics* 17, 113–142.
- Dolley, J. C., 1933. Characteristics and procedures of common stock split-ups. *Harvard Business Review* 316–332.
- Eckert, C., Gatzert, N., Schubert, M., 2023. Analyzing spillover effects from data breaches to the US (cyber) insurance industry. *European Journal of Finance* 29, 669–692.
- Fama, E. F., Fisher, L., Jensen, M. C., Roll, R., 1969. The adjustment of stock prices to new information. *International Economic Review* 10, 1–21.
- Fama, E. F., French, K. R., 1992. The cross-section of expected stock returns. *Journal of Finance* 47, 427–465.
- Florackis, C., Louca, C., Michaely, R., Weber, M., 2023. Cybersecurity risk. *Review of Financial Studies* 36, 351–407.
- Gordon, L. A., Loeb, M. P., Sohail, T., 2010. Market value of voluntary disclosures concerning information security. *Management Information Systems Quarterly* 34, 567–594.
- Gordon, L. A., Loeb, M. P., Zhou, L., 2011. The impact of information security breaches: Has there been a downward shift in costs? *Journal of Computer Security* 19, 33–56.
- Hilary, G., Segal, B., Zhang, M. H., 2016. Cyber-risk disclosure: Who cares? Available at <http://dx.doi.org/10.2139/ssrn.2852519>
- Hugues, B. B., Bohl, D., Irfan, M., Margolese-Malin, E., Solorzano, J. R., 2017. ICT/Cyber benefits and costs: Reconciling competing perspectives on the current and future balance. *Technological Forecasting and Social Change* 115, 117–130.
- Islam, M. S., Wang, T., Farah, N., Stafford, T., 2022. The spillover effect of focal firms' cybersecurity breaches on rivals and the role of the CIO: Evidence from stock trading volume. *Journal of Accounting and Public Policy* 41, 106916.

- Jamilov, R., Rey, H., Tahoun, A., 2023. The anatomy of cyber risk. Available at: <https://ssrn.com/abstract=3866338>
- Jensen, J., Paine, F., 2023. Municipal cyber risk. Available at <https://weis2023.econinfosec.org/wp-content/uploads/sites/11/2023/06/weis23-jensen.pdf>
- Jiang, H., Khanna, N., Yang, Q., Zhou, J., 2023. The cyber risk premium. *Management Science* Forthcoming.
- Johnson, M., Kang, M. J., Lawson, T., 2017. Stock price reaction to data breaches. *Journal of Finance* 16, 1–13.
- Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., Stulz, R. M., 2021. Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics* 139, 719–749.
- Karafiath, I., 1994. On the efficiency of least squares regression with security abnormal returns as the dependent variable. *Journal of Financial and Quantitative Analysis* 29, 279–300.
- Khansa, L., Cook, D. F., James, T., Bruyaka, O., 2012. Impact of hipaa provisions on the stock market value of healthcare institutions, and information security and other information technology firms. *Computers and Security* 31, 750–770.
- Kolari, J. W., Pynnonen, S., 2010. Event study testing with cross-sectional correlation of abnormal returns. *Review of Financial Studies* 23, 3996–4025.
- Kolari, J. W., Pynnonen, S., 2011. Nonparametric rank tests for event studies. *Journal of Empirical Finance* 18, 953–971.
- Kothari, S. P., Warner, J. B., 2007. Econometrics of event studies. In: Eckbo, B. E. (ed.), *Handbook of Corporate Finance*, Elsevier.
- Lee, S. H., Varela, O., 1997. An investigation of event study methodologies with clustered events and event day uncertainty. *Review of Quantitative Finance and Accounting* 8, 211–228.

- Lending, C., Minnick, K., Schorno, P. J., 2018. Corporate governance, social responsibility, and data breaches. *Financial Review* 53, 413–455.
- Liu, J., Marsh, I. W., Xiao, Y., 2022. Cybercrime and the cross-section of equity returns. Available at: <http://dx.doi.org/10.2139/ssrn.4299599>
- Malatesta, P. H., 1986. Measuring abnormal performance: The event parameter approach using joint generalized least squares. *Journal of Financial and Quantitative Analysis* 21, 27–38.
- Patell, J. M., 1976. Corporate forecasts of earnings per share and stock price behavior: Empirical test. *Journal of Accounting Research* 14, 246–276.
- Romanosky, S., 2016. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity* 2, 121–135.
- Salinger, M., 1992. Standard errors in event studies. *Journal of Financial and Quantitative Analysis* 27, 39–53.
- Savickas, R., 2003. Event induced volatility and tests for abnormal performance. *Journal of Financial Research* 26, 165–178.
- Shleifer, A., Vishny, R. W., 1997. The limits of arbitrage. *Journal of Finance* 52, 35–55.
- Spanos, G., Angelis, L., 2016. The impact of information security events to the stock market: A systematic literature review. *Computers and Security* 58, 216–229.
- Tosun, O. K., 2021. Cyber-attacks and stock market activity. *International Review of Financial Analysis* 76, 1–15.
- Zellner, A., 1962. An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association* 57, 348–368.

Appendix

Variable	Description	Source
Firm size (ln)	ln(market equity [prc*shrout])	CRSP
Firm Age (ln)	ln(years) since the firm first appeared in Compustat	Compustat
Book to market ratio	Common equity [ceq] / market equity	Compustat and CRSP
Price/Earnings	Stock Price / Earnings [pe_exi]	WRDS Financial Ratios

Table A1: **Variable definitions**

The names of the variables as found on CRSP and Compustat are in brackets.

Data availability and replication materials

To ensure transparency and replicability while avoiding excessive manuscript length, all replication materials are made available in a dedicated code repository. For anonymity reasons, the repository link is not disclosed in the manuscript and is provided separately through the journal submission system.

The repository contains two main datasets. The first dataset reports the full universe of 270 validated cyber-incident headlines identified through an NLP-based filtering procedure applied to Refinitiv news data, followed by manual verification. The second dataset corresponds to the final event-study sample used in the empirical analysis and includes the subset of events affecting publicly listed firms that satisfy all data-availability and sample-selection criteria. Inclusion in this dataset defines whether an event enters the event-study analysis.

The repository further provides all Python notebooks used for data construction, filtering, and empirical analysis. Detailed documentation describing the mapping between the full incident universe and the final empirical sample is provided in the repository.




Fig. A1: **Average abnormal returns - Zero benchmark**

Average abnormal daily returns in an 11-day window centered on the cyber events. Day 0 is the day of the cyber event. The OLS limited model is the OLS model restricted to the same period, the same subset of firms and events as the SUR model. Abnormal returns are computed with respect to the zero-benchmark model, following Eq. 2.




Fig. A2: **Cumulative average abnormal returns - Zero benchmark**

Cumulative average abnormal daily returns in an 11-day window centered on the cyber events. Day 0 is the day of the cyber event. The OLS limited model is the OLS model restricted to the same period, the same subset of firms and events as the SUR model. Abnormal returns are computed with respect to the zero-benchmark model, following Eq. 2.